Planning chemical syntheses with deep neural networks and symbolic AI

Article presentation by Benjamin Heuser,
Frank de Morrée and Iestyn Watkin.

Nature (29 March 2018)
doi:10.1038/nature25978

Outline

The Problem

Computer-aided retrosynthesis is slow and unsatisfactory.

- ➔ Retrosynthetic analysis is the general technique used by chemists
- ➔ Used to plan the synthesis of small organic molecules

Retrosynthesis

"In retrosynthesis, a search tree is built by 'working backwards', analysing molecules recursively and transforming them into simpler precursors until one obtains a set of known or commercially available building-block molecules."

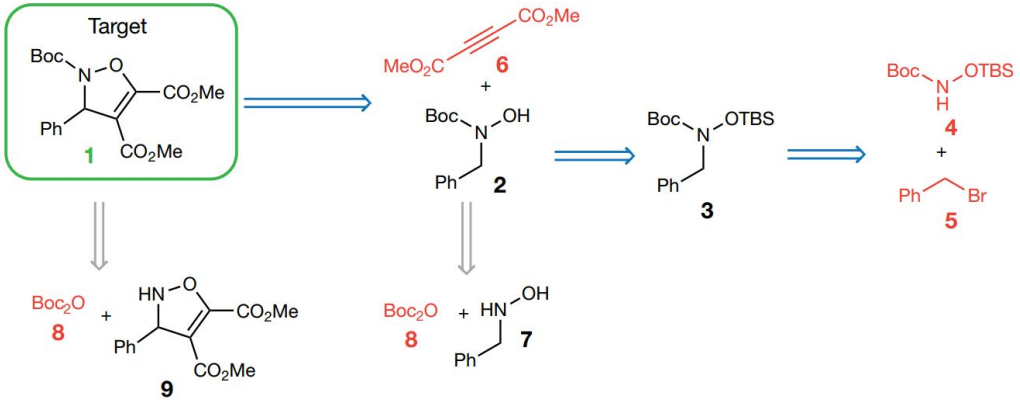

Methods

The 3N-MCTS method.

Three neural networks in combination with a Monte Carlo search tree method were employed to solve the problem.

Data selection

A.

The data consisted of two subsets to train the neural networks.

Data = transformation rules

- ➔ Taken from Reaxys chemistry database listing all known reactions
- ➔ Indication: 17,134 and 301,671 rules kept for the neural networks.

Monte Carlo tree search

B.

Explores a large amount of options making the Monte Carlo method is a good choice.

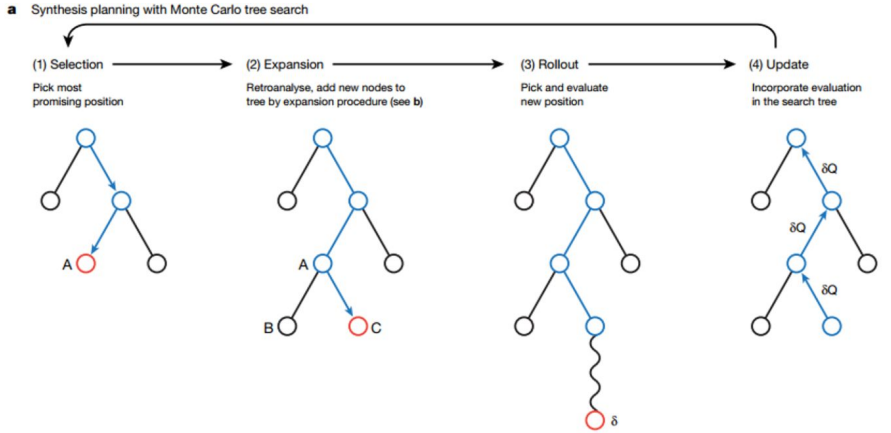

Figure 1: Schematic of MCTS methodology. **a**, MCTS searches by iterating over four phases. In the selection phase (1), the most urgent node for analysis is chosen on the basis of the current position values. In phase (2) this node may be expanded by processing the molecules of the position A with the expansion procedure, which leads to new positions B and C, which are added to the tree. Then, the most promising new position is chosen, and a rollout phase (3) is performed by randomly sampling transformations from the rollout policy until all molecules are solved or a certain depth is exceeded. In the update phase (4), the position values are updated in the current branch to reflect the result of the rollout.

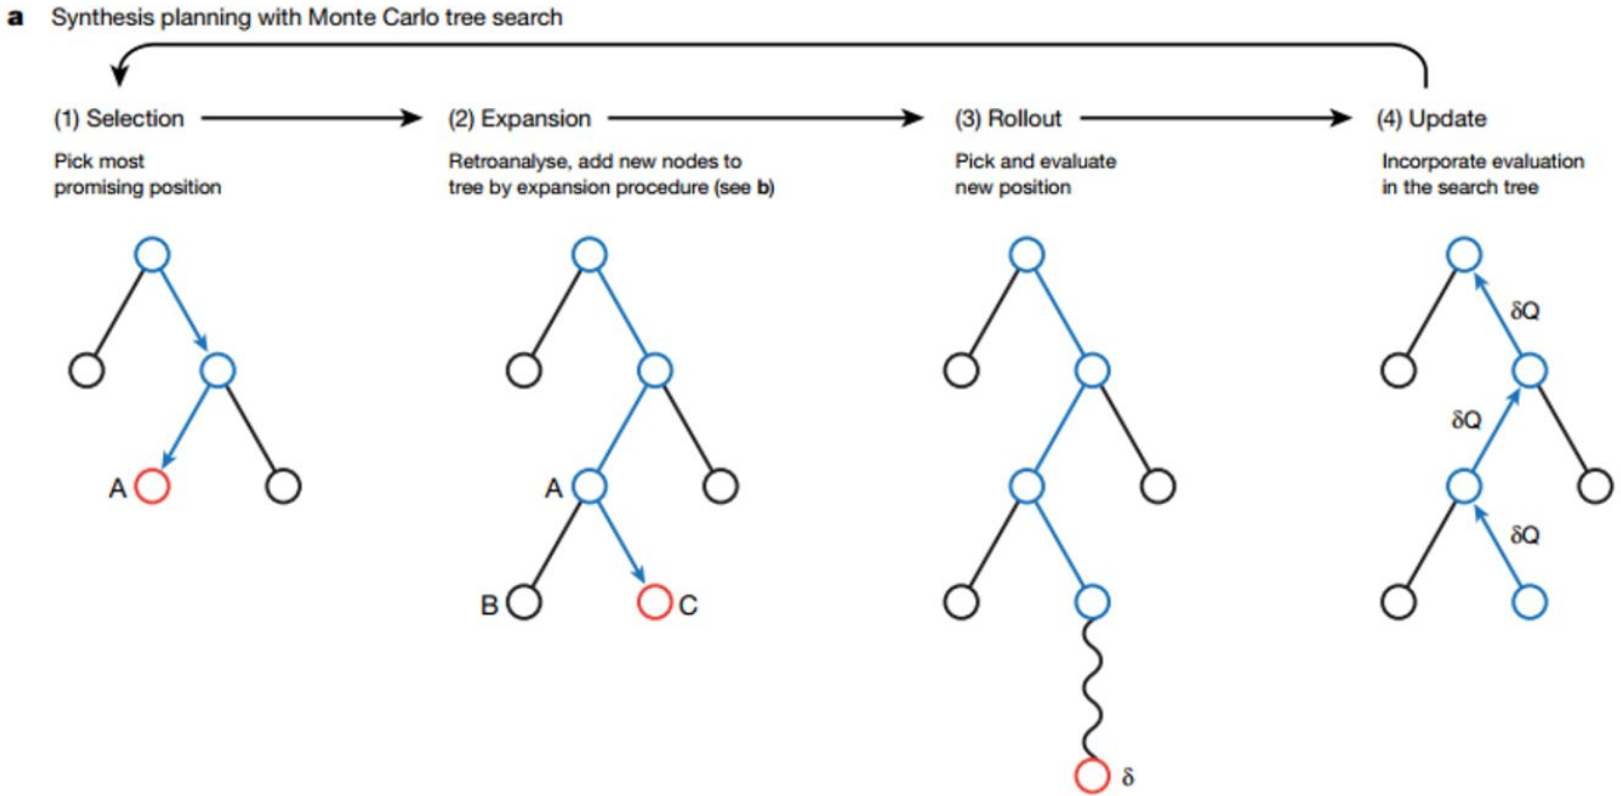

Figure 1: Schematic of MCTS methodology. **a**, MCTS searches by iterating over four phases. In the selection phase (1), the most urgent node for analysis is chosen on the basis of the current position values. In phase (2) this node may be expanded by processing the molecules of the position A with the expansion procedure, which leads to new positions B and C, which are added to the tree. Then, the most promising new position is chosen, and a rollout phase (3) is performed by randomly sampling transformations from the rollout policy until all molecules are solved or a certain depth is exceeded. In the update phase (4), the position values are updated in the current branch to reflect the result of the rollout.

Selection phase

C.

Does something.

$$a_t = \underset{a \in \mathcal{A}(s_t)}{\operatorname{argmax}} \left(\frac{Q(s_t, a)}{N(s_t, a)} + cP(s_t, a) \frac{N(s_{t-1}, a_{t-1})}{1 + N(s_t, a)} \right) \tag{2.1}$$

1. $Q(s_t,a)$ is the action value.
2. $N(s_t,a)$ is visit count to current node.
3. $P(s_t,a)$ is the prior probability
4. $N(s_{t-1},a_{t-1})$ is the state-action pair that led to the current node.

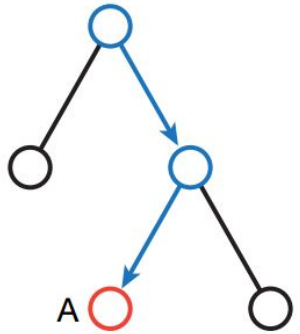

Expansion phase

D.

Expands previously visited leaf nodes.

1. Highway neural network (allows shortcuts)
2. Trained on chemical synthesis data to predict the most probable transformation with its probability
3. Suggest the 50 most likely transformation and adds to tree
4. Is only triggered once if the leaf node is visited for the second time
5. New leaf nodes are yet unrated

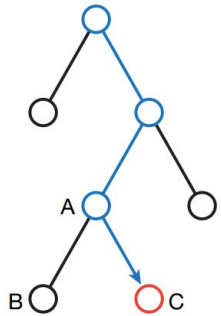

In-scope filter E.

Filters unfeasible reactions.

1. Neural network trained to predict whether a proposed reaction will work (binary classifier)
2. Trained on reaction data for single step reactions <= 3 reactants and 1 product
3. Negative training data artificial created by perturbation

Rollout phase

F.

Determines the position value of a node.

1. Neural network trained to sample the child nodes
2. Triggered if leaf node visited for the first time
3. Recursively samples child nodes until building block is reached or recursion depth
4. Casts a rewards function based on solved states

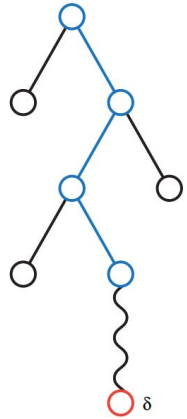

Update phase
G.

Backpropagate the rewards through the tree

1. The rewards values are added to the parent tree nodes
2. Long paths are punished
3. Paths that are longer than a maximum depth of 25 edges are removed
4. The visit count of the nodes are updated to keep track of how often each node has been visited

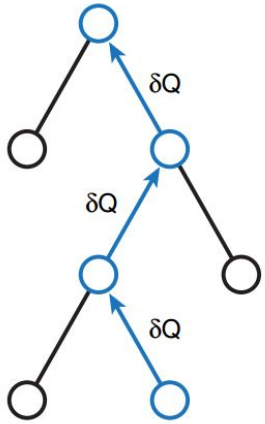

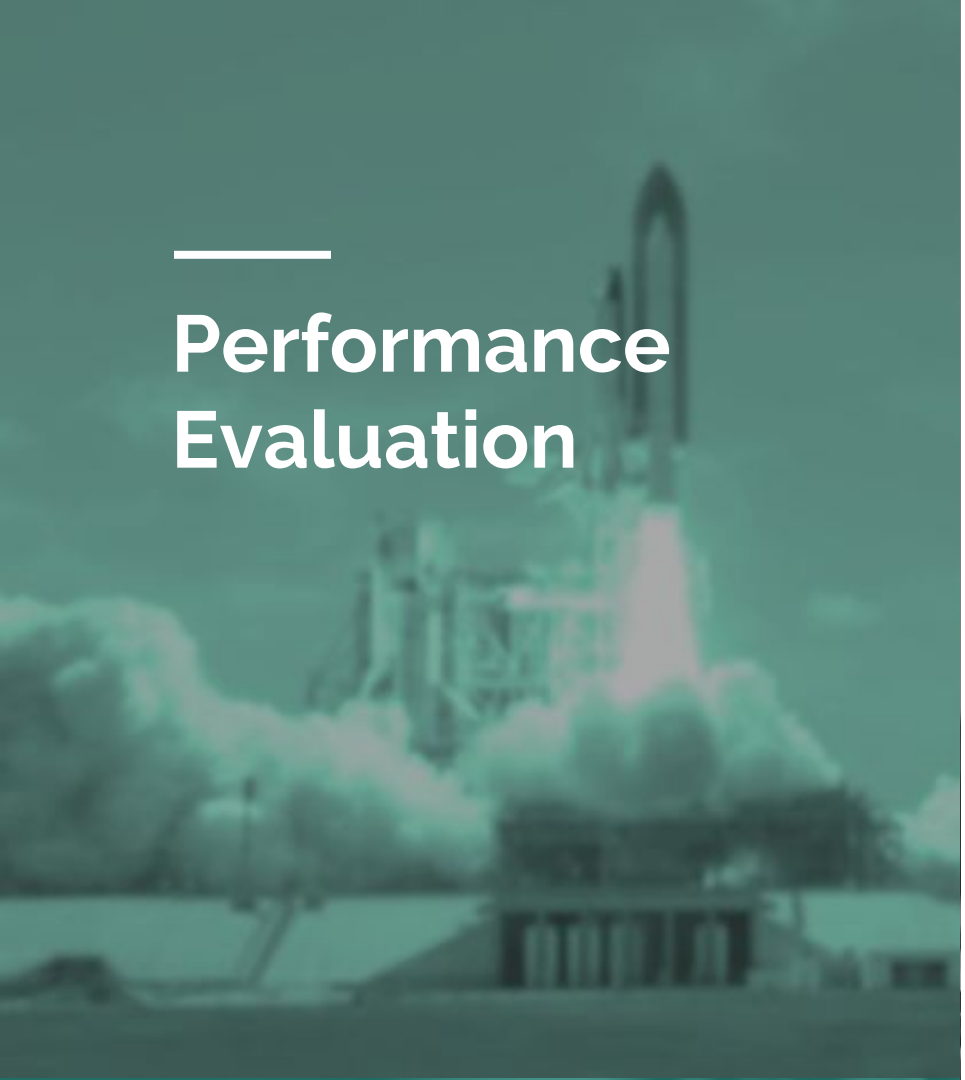

Performance Evaluation

So how did we do?

The 3N-MCTS method was compared to other modern methods and it’s reaction routes judged by expert chemists.

Results

- ➔ Expansion policy network 31% correct out of 301,671 reaction transformations: **Reasonable!**
- ➔ Accuracies of 63.3% and 72.5% for top 10 and top 50 results for the expansion network: **Good!**
- ➔ In-scope filter achieved area under ROC values of 0.99 on the test set: **Good performance!**
- ➔ Compared to other methods, 3N-MCTS was found **30x faster** and solves for **2x more molecules**.
- ➔ Wilcoxon signed rank test: 45 graduate-level chemists chose reactions, found machine routes not to be inferior.

Discussion and Conclusions

What did we learn and what's next?

There are still some difficulties but the performance upgrade is good.

Computer-aided Retrosynthesis

Using Monte Carlo tree search in combination with 3 neural networks, retrosynthetic routes were found. The performance is higher than canonical methods, yet more research is required.

Difficulties

- ➔ Sparsity of training data
- ➔ Natural product synthesis impossible
- ➔ More advanced methods required to aid humans better

3N-MCTS conclusions

- ➔ 30x speedup
- ➔ 2x larger solution space
- ➔ Adding complexity from neural networks is valuable
- ➔ Expert chemists find machine generated routes of good quality

References

M. H. S. Segler, M. Preuss & M.P. Waller, “Planning chemical syntheses with deep neural networks and symbolic AI”, Nature 555, 604-610 (29 March 2018).